



# Dropbox Integrated AV Production Environment

David Carroll - Founder & President  
Kevin Gross - Media Network Architect



**David Carroll**  
Associates, Inc.  
Your Technology Integrator

## Table Of Contents

About David Carroll Associates, Inc.	3
Published By:	3
Introduction	4
Requirements and Functional Design	4
Converged Networking	5
Network Topology	7
PTP Topology	11
Audio Systems	11
Video Systems	12
Lessons Learned	13
PTP Fault Tolerance	13
PTP Multicast Delivery	15
PTP Disruptions	15
PTP Capacity	16
Multicast Routing	17
Switch Software	18
IP Addressing and Routing Tables	18
Monitoring and Management	18
Conclusion	19
Sidebar: Audio Protocols	20
AES67	20

# Dropbox Integrated AV Production Environment

Ravenna	20
Q-LAN	21
Dante	21
VoIP and RTP	21
<b>Sidebar: Synchronization Protocols</b>	<b>22</b>
Precision Time Protocol	22
Network Time Protocol	23

## About David Carroll Associates, Inc.

David Carroll Associates sets the standard for AV/Technology Systems Design and Integration. The company melds its rich background in broadcast, film, production and live sound projects with intellectual enthusiasm and relentless passion.

DCA specializes in designing, installing and maintaining integrated audio, video and control systems for a wide variety of environments, including educational settings, medical facilities, corporate environments, auditoriums, teleconferencing settings and more.

Every client gains the benefit of DCA's deep experience and insight as the engineers dig deep to understand its unique goals and working environment. DCA applies its unique "engineering approach" to every project: understanding the client's requirements, making design choices that best meet the client's needs, and apply traditional engineering disciplines (such as documentation) to make sure everything works cohesively. The results are all custom solutions — individually tailored, impeccably installed and poised for the client's next phase of growth.

## Published By:

David Carroll Associates, Inc.

1320 S. 51st St.

Richmond, CA 94804

Phone: (510) 235-1256

[info@davidcarroll.com](mailto:info@davidcarroll.com)

[www.DavidCarroll.com](http://www.DavidCarroll.com)

© 2020 David Carroll Associates, Inc. All Rights Reserved.

May be shared with copyright and credit intact.

## Introduction

New interoperable network protocols eliminate most non-network interconnects and provide fault tolerance based on IT standards and techniques. The networked approach gives users flexibility for on-the-fly production from anywhere in the facility. This network-based approach is particularly suitable to network-centric technology companies.

Compared to broadcast applications where many of the early rollouts of SMPTE ST 2110 have occurred, corporate and commercial applications of networked audio and video technology have different requirements. These differences result in different network design and system buildout.

A key network design decision for any AV network is network topology and relates to the scale and physical layout of AV equipment. A leaf-spine topology is commonly used for media networking. Different topologies tend to be used for enterprise networking. A hybrid topology can serve multiple needs in a converged network.

The converged network combines media and control on the same media network and integrates this media network with the larger enterprise network and provides additional flexibility while making the system easier to manage and potentially less complex and costly. Achieving this promise requires a skilled integrator with strong networking knowledge and experience.

Given the state-of-the-art approach, the hybrid topology used for networking and PTP, there were a number of challenges to overcome in designing and integrating the system. Assuring fault tolerance for critical PTP synchronization and reliable delivery of accurate synchronization to different types of devices at different levels of the network topology without overloading any of the equipment.

## Requirements and Functional Design

Dropbox has recently commissioned a new corporate headquarters. Internet-focused companies such as Dropbox see the internet as a new form of media and strive to apply internet techniques and technology to traditional media production. To this end, Dropbox included embedded AV systems in the buildout of its headquarters. The new network-based production infrastructure is particularly attractive for these clients as it represents a convergence of AV and network technologies, a general trend many of these companies have based their businesses on.

The networked approach means that, in principle, all signals are available anywhere and anything can be controlled from any point. The types of media carried over the network include AES67, Q-Sys Q-LAN audio, Ravenna, and Dante audio; compressed video in various forms; SMPTE ST 2110 audio and video streams. The network also carries control and monitoring protocols of various types and supports event logging and system monitoring diagnostic systems.

# Dropbox Integrated AV Production Environment

Network fault tolerance standards for general IT in the organization were already well established. These standards allow for single points of failure towards the edge of the network but require redundancy towards the core. Failure of a network edge device may cause loss of connection to the directly connected devices, but the failure of a central network device does not affect connectivity for end stations.

A significant thrust of the installation was providing flexibility for operators to be creative in how media is produced. The system provides decentralized production capabilities throughout the facility. At Dropbox, the main areas of media interaction include the dining facility used for meals and corporate social events and a series of classrooms with scalable teleconferencing capabilities. A central control room provides centralized operations and production support for many of the productions, but production can also be controlled locally where it is happening. To this end, over 80 broadcast service point (BSP) wall plates are available throughout the facility. The BSP provides a fiber-optic connection for simple handheld video cameras to more sophisticated production carts or "briefcases" and provide several levels of technical capabilities for covering these in-situ events in spontaneous and creative ways.

## Converged Networking

A prevailing approach to professional media networking in production facilities is to install a dedicated network for media equipment. In applications requiring fault tolerance, two separate and identical media networks may be built. Many media endpoints feature multiple network interfaces: one or two for real-time media data plus additional connections for control and monitoring. A separate control network is often installed in these facilities to keep real-time media physically separated from control and monitoring and other traffic.

This approach is an improvement compared to how non-networked systems have been built with central audio and video routers connected directly to each media endpoint. The network approach uses standard IT equipment and interconnects and offers endpoints flexibility in the format of the media carried by the network.

Further improvement and new capabilities become available when the audio, video and control networks are more integrated and, for organizations where professional media production is not the central business focus, integration with the business IT infrastructure can be especially fruitful. This is known as a converged approach and refers to the convergence of AV and IT functions on a single network. The additional opportunities for connections on a converged network introduce new flexibility in how the media systems can be used and media systems can be more readily integrated with the business systems creating synergy for both.

The converged network represents a more IT-like approach to networking. The mission of enterprise IT is to build a network that seamlessly serves the needs of users. IT does not view this mission as building individual networks for individual needs but supporting a single network with the capabilities and capacity to satisfy all needs. The flexibility, connections, and scalability of these enterprise networks are what produce the business benefits.

# Dropbox Integrated AV Production Environment

Convergence requires building a network to meet the requirements of the most demanding application. Although this may seem like an expensive proposition, it is already necessary to meet those requirements for at least part of your system and extending these capabilities to other areas does not represent a huge incremental cost. A more capable network serving all your applications is a valuable upgrade to an increasingly critical part of the AV system.


Implementation choices in some products and the daunting chore of assuring a network-connected product will play well with hundreds of other possible network devices and network equipment from different vendors cause manufacturers to recommend their equipment be installed with network equipment with which they have direct experience and to be separated from other network devices with which they are not required to directly interoperate. Fully accommodating these manufacturer recommendations means backing away from convergence and building multiple networks or multiple virtual networks on a single infrastructure. This may be more costly, is less flexible and introduces overhead in network configuration, monitoring, and maintenance activities.

This makes sense from the manufacturer's perspective, but it doesn't necessarily serve the needs of the installation. The job of an integrator is to make things work together. In a networked world, that means making things work together on the same network. An integrator with strong network expertise is one way to bridge the gap between the level of network sophistication manufactures can support and what a state-of-the-art installation aspires to.

If convergence is achieved, the number of devices on the same network can become difficult to manage. For large projects such as Dropbox headquarters, naively putting everything on a single network may exceed the scale of what Ethernet is able to reliably support. A best practice is to partition the network into local subnetworks and use the Internet Protocol to build a scalable network of networks. Having a subnetwork or two associated with an equipment room will partition problems occurring in that location from affecting other locations. This is in contrast to the un converged approach of having a network per application and configuring or building that network to span all equipment rooms. In this case, if there is a problem with one of the applications in one place, it affects the entire facility and is difficult to determine from where the problem originates.


## Network Topology

Modern network topologies tend to be hierarchical because hierarchy is key to scalability. Scalability is the adaptability of a system to different sized applications. In the case of networking, scales range from the home network, to the office, to the enterprise to the entire internet.






# Dropbox Integrated AV Production Environment




The smallest and simplest network is a central switch to which all endstations are attached. A two-tier hierarchical system moves endstations to edge switches that are connected to a core switch. A three-tier system adds an internal distribution layer to address the limit to the number of connections that can be concentrated at the core. In media and datacenter applications it is popular to refer to a two-tier topology as *leaf-spine*. In this terminology, the spine is the core and leaves are edges.



# Dropbox Integrated AV Production Environment

Through this range of hierarchical networks, there is the option to add additional equipment to the network to achieve fault tolerance. A simple example in the simplest network is the addition of a second central switch with all endstations connected to both switches. If one of the switches or if any connections fail, there is still a working connection path for endstations. The doubling-up of equipment and connections can be extended to the larger topologies.

In a pure hierarchical topology, all endstations attach to the edge or leaf switches at the bottom layer of the hierarchy. Hybrid hierarchical topologies allow the connection of endstations at other levels of the hierarchy. So, for instance, demanding endstations such as those supporting uncompressed video may be connected at the core level where bandwidth requirements may be more readily satisfied. Control and monitoring endstations remain connected at the lowest level.




At Dropbox, we chose to use a two-tier hybrid hierarchical topology with redundant core switches for fault tolerance. The SMPTE ST 2110 video equipment is attached directly to the core switch in all cases and all other equipment, including stand-alone audio and all control and monitoring functions, are connected at the bottom layer of the hierarchy.

We used Arista chassis switches at the core and Cisco stackable switches at the edge. The uplinks between the core and edge are 10 gigabit Ethernet: a pair of single-mode fiber from each edge location to each of the core switches.

# Dropbox Integrated AV Production Environment

Professional media protocols including Dante, Ravenna, Q-LAN, and SMPTE ST 2022-7 support redundant streaming. A sender supporting redundant streaming transmits two copies of media stream data, typically from two different network connections. These are often labeled, for instance, "red" and "blue" to distinguish the two identical sides of redundant streaming. Receivers are similarly equipped and receive two copies of data to work with. If one copy is corrupted, missing, or arrives late, the other good data is used, and playback is uninterrupted. Redundant streaming protects against intermittent data loss as well as more catastrophic failure scenarios.

Beyond the "red" and "blue" networks supporting professional video, other sections of the network are subdivided based on function and location. IP routing is used to link the separate areas together. Putting separate systems on separate IP subnets creates a degree of fault isolation for the physical and functional areas of the network.



In a continuation of the IT-centric design, the media network does not stand completely apart from the enterprise network. Media equipment occupies the same private IP address space as other enterprise network users. A set of two 10 gigabit uplink connections between the core switches from the media network connect it to the corporate network. The IT-standard Open Shortest Path First (OSPF) routing protocol is used to exchange routing table information to allow communications with corporate facilities.

# Dropbox Integrated AV Production Environment

Secure internet access is also provided through the enterprise connection as access to the internet for updates and licensing is an increasingly common requirement for production equipment and the software that supports it.

It may be undesirable for users of the media network to have access to sensitive corporate information. Conversely, it may be desirable to restrict access to sensitive media to employees at large. Access control lists (ACLs) are configured in network equipment to restrict access between the media network and the larger enterprise network.

## PTP Topology

Synchronization is a critical service on professional media Networks. The source for synchronization is typically a master clock generator. The master clock is typically synchronized to a traceable source such as GPS. Precision Time Protocol conveys synchronization information from the master clock to devices requiring synchronization.

PTP defines a hierarchical topology separate from the network topology for synchronization distribution. In PTP parlance, the master clock generator is known as the *grandmaster*. The simplest PTP topology is for all devices requiring synchronization to synchronize directly to the grandmaster. The size of a network using this topology is limited by the grandmaster's capacity. To address a potential bottleneck at the grandmaster, PTP hierarchy includes the concept of a *boundary clock*, a secondary clock synchronized to the grandmaster, and providing synchronization signals to local devices.

The Dropbox network uses a two-tier PTP topology. Two Tektronix SPG-8000A master clocks receive a GPS time reference. Using PTP's best master clock algorithm (BMCA) one of the two is selected to deliver PTP synchronization to the boundary clocks in the Arista core switches. The boundary clocks provide synchronization directly to critical video devices. PTP synchronization is provided indirectly by the boundary clocks to devices connected to the Cisco edge switches.

## Audio Systems

The facility audio systems serve background music and sound reinforcement for meeting and classroom activities. The classrooms also have a teleconferencing capability including acoustic echo cancellation.

The core of the facility's audio system is the QSC Q-Sys system. Q-Sys supports several audio networking protocols including its native Q-LAN protocol, AES67, Dante, and several grades of voice over IP (VoIP). The supported protocols allow connections to many audio sources such as network-powered ceiling speakers, wireless microphones and other room microphones, and networked amplifiers for sound reinforcement and background music.

# Dropbox Integrated AV Production Environment

The facility audio system exchanges networked audio signals with production systems. The core of the production audio system is a Lawo V-Matrix system dedicated to audio routing and two Lawo MC<sup>2</sup>36 consoles. These core components of the production audio systems use PTP synchronization and are connected directly to the Arista core switches. The Lawo system uses AES67, Ravenna, and 2110 audio protocols.

Facility audio sources such as wireless and fixed microphones for sound reinforcement and teleconferencing are also used as primary audio to the production systems. Conversely, production system playout may be routed to the facility sound system to support theater and meeting overflow scenarios.

This blurring between facility audio and the production is possible because two systems share a common media network and are synchronized to a common clock. Aside from the operational scenarios identified during system design, the ability to interconnect systems through a common network is expected to produce additional flexibility and creative benefits throughout the life of the system.

## Video Systems

Video comes into the system either as native 2110 or as 3G SDI. Lawo V-Matrix systems along with the Arista switches form the core of the video system. While the Arista switches are centralized, the V-Matrix frames are located throughout the facility wherever video is used in abundance, specifically in machine rooms associated with classrooms and auditoriums.

Each V-Matrix frame can accommodate up to 8 C100 video processing cards. Each C100 card has SDI inputs and outputs and two 40 Gb optical Ethernet port capable of supporting multiple 2110 streams. These cards perform functions from simple conversion of audio and video between 2110 and SDI to multiviewer support.

Video sources, destinations and processing lie outside the core video infrastructure and include SDI and 2110 native video cameras; Grass Valley 2110 native K-Frame switcher; Evertz Dreamcatcher multichannel native 2110 video record and playback; Tektronix Prism endpoint analyzer; AWS Elemental Live encoders for video distribution over the company intranet and Aja Ki Pro storage for conventional offline production.

The K-Frame switcher is partitioned into two suites, the larger one for conventional centralized production and a smaller 16 x 8 suite for production operations where an event is happening using one of the BSP connections located throughout the facility. The BSPs can be dynamically configured to support cameras directly, support connection of native 2110 devices, and also support more general networking for audio streaming and remote control and monitoring.


## Lessons Learned

### PTP Fault Tolerance

Professional AV systems typically cannot operate without a reference clock. In a 2110 system, the reference clock comes from the PTP grandmaster. The grandmaster is a potential single point of failure for the system, so it makes sense to design some redundancy into this area.

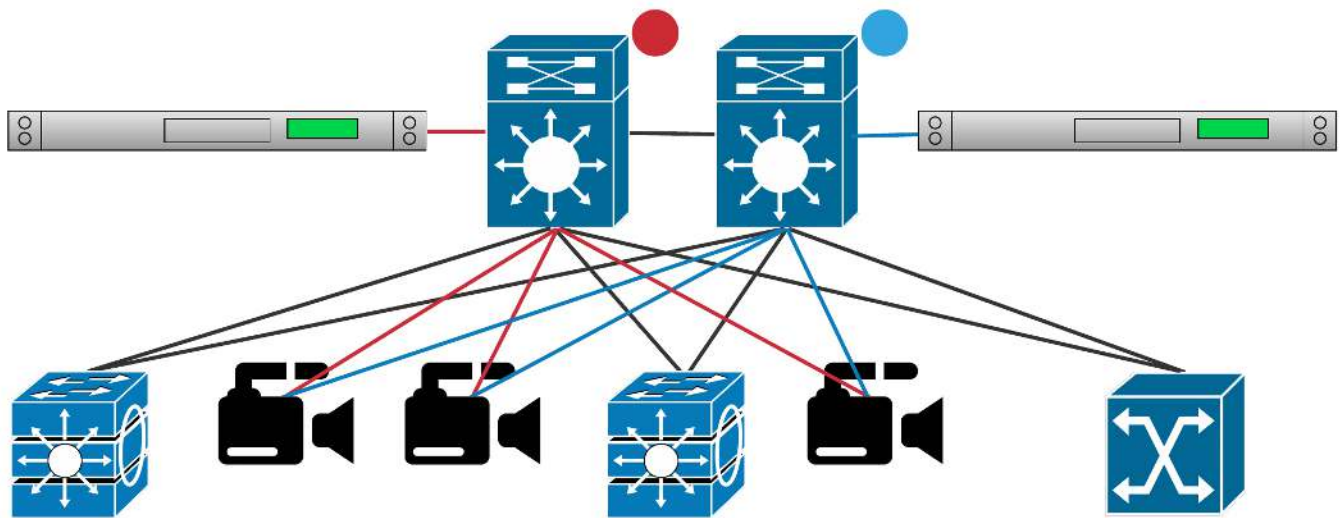
In networks built to support 2022-7 redundant streaming with two separate isolated networks, a designer might be tempted to put a grandmaster on each network. With both grandmasters referenced to GPS, these two clocks can theoretically be used interchangeably, and it can be left to connected devices to determine which grandmaster to synchronize to.

This simple approach starts to fall apart as failure scenarios are considered. For example, if one or both of the grandmasters lose their GPS connection, the clocks will begin to drift and can no longer be used interchangeably. If one of the grandmasters fails completely, the associated boundary clock will assume grandmaster duties based on its own drifting internal clock.



A more robust approach to this problem is to allow the two grandmasters to interact through the BMCA and self-select which grandmaster should be active. With isolated networks, this can be achieved by adding a special PT-only link between the networks.

# Dropbox Integrated AV Production Environment



Because the Dropbox network uses a converged approach and features routed IP connections between the "red" and "blue" networks, the two grandmasters can be directly connected to the cores and we avoid the expense, complexity, and single point of failure associated with the additional PTP hierarchy.

In addition to grandmaster fault tolerance, it is necessary to consider fault tolerance in PTP boundary clocks and the paths they use to distribute synchronization. Once the grandmaster is selected, the boundary clock in the core switch it is connected to receives synchronization from the grandmaster directly and the other boundary receives synchronization from the first.

All PTP elements are receiving a clock and the synchronization system is working at this point, but, as with most complex systems, it can be valuable to examine system state more closely and verify that everything is working *as intended*.

On closer inspection, it turns out there are multiple paths between the two boundary clocks. The path we prefer is the direct connection between cores. The other paths are out to any edge switch at back. Since we haven't done anything to tell the network our preferred path, one of these edge-switch paths is chosen and the second core boundary clock receives a workable but degraded clock signal. It turns out that Arista switches do not have a means of configuring PTP path priority but make this selection based on port number. To assure our direct core-to-core connection was the preferred PTP path, it was necessary to move this connection to a low-numbered port on the Arista core switches.

## PTP Multicast Delivery

PTP uses IP multicast messaging to communicate synchronization. On conventional networks, this messaging is carried through normal IP multicast routing mechanisms. When PTP-aware network equipment is used, these messages are intercepted, interpreted, and potentially retransmitted outside of normal IP multicast routing mechanisms.

Some portions of the Dropbox AV network use PTP-aware equipment while others do not. It is crucial to manage the boundary between these two modes of operation and to properly configure ports on each side of the boundary. Failure to handle this interface properly can cause problems in PTP message delivery through the boundary resulting in failure to synchronize and disruption of protocols supporting the network's IP multicast routing.

## PTP Disruptions

Accurate synchronization relies on reliable and prompt delivery of PTP messages across the network. Unfortunately, packet loss and delay is common in networks. Because packet loss and delay is commonly experienced, when possible, network protocols are designed to be tolerant to this. The sensitivity of PTP leads to situations where network engineers have evidence that the network is operating properly based on experiments with more tolerant protocols but there are, in fact, lower level problems with the network adversely affecting media performance. These problems may first become apparent in media applications and specifically PTP synchronization.

Specific packet loss issues affected PTP unacceptably in the Dropbox installation but had a negligible effect on control communications and caused only minor issues for media data. These issues were tracked down to internal switch issues that required a software update to resolve and marginal fiber connections that required cleaning or retermination.

The network must be properly engineered and configured to avoid mistreatment of PTP messages. Defects in network equipment or connections have the potential to cause packet loss. In PTP-aware equipment, these messages are automatically prioritized on ports where PTP awareness is enabled. In other portions of the network DiffServ QoS is used to identify and prioritize PTP messages. Failure of either of these mechanisms to do their job can cause PTP disruptions.



# Dropbox Integrated AV Production Environment

If the *Sync* messages from the master clock are lost, a slave device loses opportunities to adjust its local clock and will drift away from ideal synchronization. If messages are lost or delayed in the *Delay\_request/response* conversation, the inability to update the reading for the network delay from master to slave can cause synchronization inaccuracy. A good PTP implementation can typically tolerate these disruptions. More challenging is the case where these messages eventually arrive but have been delayed by the network. An unsophisticated PTP implementation may make wild adjustments based on delayed packets and lose synchronization. Even a robust PTP implementation is likely to deviate from ideal synchronization in the presence of these network disruptions. The effect of the deviations depends on the synchronization accuracy requirements of connected equipment.

On the Dropbox, our most demanding PTP slave are the antennas serving the Riedel Bolero wireless intercom system. These intercoms use the DECT radio protocol which requires participants to accurately time their transmissions based on PTP time to avoid interfering with one another in the radio spectrum. A poor PTP clock is detected by these devices and intercom audio is muted in both directions until clock quality recovers.

A more serious type of disruption is indecision about the best master clock in the BMCA. The BMCA relies on regular *Announce* messages from the current and potential master clocks. If the network does not reliably deliver *Announce* messages from the grandmaster, other potential masters may invoke the BCMA and assume the grandmaster role. There is typically a disruption for slaves as they synchronize to a new grandmaster. Also, in this scenario, due to the communication issues, there are likely more than one grandmaster operating, at least briefly, on the network. In order to successfully exchange media, devices must be synchronized to the same grandmaster. With multiple grandmasters, some media connections will be disrupted. The use of the *Slave-only* PTP configuration option advocated in 2110 can reduce or eliminate the possibility of multiple grandmasters.

## PTP Capacity

PTP, being a master-slave protocol, puts a performance burden on the grandmaster or boundary clocks and the network connections leading there. The process in the clock of receiving and responding to requests from slaves and sending regular time announcements is handled by a CPU. In the case of PTP-aware network equipment, the CPU is also responsible for other high-level networking functions. PTP capacity is therefore limited by CPU capacity and other demands on the CPU. The effort required to support a single slave is dependent on the behavior of the slave as determined by the PTP profile in use.

It is important to understand the capacity limits of your grandmaster and PTP-aware network equipment for the PTP profile in use. Because of the number of design variables in play here, it is prudent to leave ample capacity headroom when designing a PTP synchronization system.

With PTP support in switches becoming more common, it is possible to build a network where every PTP slave device is directly connected to a switch containing a boundary clock. With PTP support in all network equipment, the burden of supporting PTP slaves is distributed and PTP capacity grows as the network grows.

## Multicast Routing

All 2110 media is transmitted using IP multicast. IP multicast is also critical to PTP synchronization and for certain connection management protocols that use IP multicast as a means of discovery. All of this makes IP multicast a crucial technology for AV networks. IP multicast is well supported in enterprise-grade network equipment. It is, however, typically disabled by default and typically remains off in enterprise deployments. Network engineers may have limited experience configuring and debugging IP multicast.

IP multicast relies on three non-trivial protocols: IGMP, PIM and whatever IP routing protocol is in use on your network (e.g. OSPF, EIGRP). Getting all this setup and avoiding undesirable interactions can be an undertaking. Testing and debugging problems with this system can be even more difficult.

In a single-vendor installation, it is possible to rely on examples from the vendor to guide configuration and you may be able to rely on technical support from the vendor to help diagnose and resolve any initial issues. Having access to a network engineer with an understanding of IP multicast is critical when dealing with multi-vendor installations and more intricate problems that may arise in more demanding production scenarios.

Multicast protocols must be designed to work with multicast routing. This includes using destination multicast IP addresses in the range that may be routed. The mDNS discovery protocol used by Dante and, to extent, NMOS uses a non-routable address. IP packets must use an appropriate TTL IP header field value for routing. Some PTP implementations including Dante's assume they will be operating only on a local area network and set their TTL field accordingly.

Some systems are able to avoid the bulk of IP multicast issues by putting all equipment on the same IP subnetwork. With everything on the same network, routing, unicast or multicast, does not come into play. This approach is only applicable to a certain scale. Avoiding IP routing also deprives you of the many scalability and robustness benefits gained from it.

## Switch Software

Modern network and AV equipment are software-intensive systems. Software tends to be a weakness in these systems. Updating network equipment software can require a scheduled maintenance window and disrupt production. It is therefore important to choose a version of software for your equipment, apply it consistently to your equipment, and test it thoroughly. Network and AV software updates may be required to resolve problems. These updates should be addressed with patience and care to avoid or understand any potential regressions. You don't want to fix one thing and end up breaking something else worse. A formal change management protocol with peer review is crucial to the maintenance of production systems.

## IP Addressing and Routing Tables

Today's AV network installations typically use static IP addressing. This is contrary to trends in IT at large where dynamic IP assignment using DHCP is frequently used. Both static and dynamic addressing require significant network engineering to partition the address space into subnetworks. This can be particularly challenging when integrating with an enterprise network.

All facilities supporting AV networking will also have an IT network for more mundane computing chores. Even if there is no planned connection between the AV network and the IT network it makes sense to coordinate IP addressing schemes to avoid any address space overlap between networks. AV engineers will need to work with the IT department to find a block of IP addresses for the AV network that doesn't overlap with IT functions. Coordinating IP address assignments with IT will simplify troubleshooting and will allow for more flexibility going forward.

If changes are needed to the subnetwork layout in a network using static addressing, it is necessary to visit each device and manually adjust its addressing configuration. The measure-twice-cut-once wisdom applies when engineering a statically addressed AV network. Be sure to leave room in your address maps for additional devices and for commissioning new subnetworks for unforeseen applications.

## Monitoring and Management

The network, as the common infrastructure for the audio and video systems, has a lot going on at all times. Tools and techniques for detecting and resolving problems and undesirable interactions are needed. Standard network diagnostic tools such as SNMP or at the CLI of network equipment provide detailed information that may be difficult or time-consuming to interpret. Because of the amount of data associated with media flows, packet captures become fire hoses and can overwhelm equipment performing the capture or the ability of applications and engineers to analyze the results. These tools can still be effective if used carefully.

# Dropbox Integrated AV Production Environment

Alternatives or supplements deployed in the Dropbox network are purpose-built tools such as the Tektronix Prism and Lawo SmartDash and SmartScope. The Prism acts as a 2110 endpoint and can do digital video analysis familiar to broadcast engineers. Since it is a native 2110 device, it also analyzes PTP behavior and any potential QoS issues with arriving data. This network-level performance information is also available in SmartScope which is essentially a 2110-aware packet capture appliance with a 100 Gbit/s network connection.

The Prism can effectively be used as a reference receiver. If a device is having trouble synchronizing or receiving a video stream, the Prism can be configured to receive the same stream to help determine whether the problem is with the transmitter, receiver, or network.

SmartDash was brought into the project after too many eyes glazed over looking at character-based information displays. The product communicates with network and video equipment and summarizes system status on a dashboard screen. This has the benefit of alerting system operators to potential problems with AV operations before the problems are reported by users and staff.

## Conclusion

AV networking promises and delivers the flexibility that enables and even promotes creativity in media production. By relying on standard IT equipment, this flexibility is achieved with reasonable cost and engineering effort. Network architecture, design and set up for a network supporting professional AV deviates significantly from standard IT practices. The most effective designs, using a convergent approach, harmonize AV, and IT design practices.

AV equipment can be significantly more sensitive to subtle network issues than other network applications due to the demanding nature of AV applications and implementation issues in individual products.

Multicast routing is critical for several aspects of AV networking. PTP uses multicast to deliver time information to multiple destinations simultaneously. 2110 uses multicast for all video stream data. AES67 also has a popular multicast option. Multicast delivery on larger networks such as these involves a number of moving parts and protocols.

Standard tools and special expertise can be used to diagnose and correct problems. Special tools may help assess network performance and quickly locate problems. Some of the newer tools can give operators ongoing status information for the entire system.

## Sidebar: Audio Protocols

Some of the early audio network protocols work only over the type of network they were designed for. For example, CobraNet only works on a simple Ethernet network, AVB requires network equipment supporting the AVB protocols, and AES47 was designed for now obsolete ATM networks.

Most of the current generation of audio protocols are IP based. IP protocols are generally more versatile and work over any IP network with suitable performance. Most of these high-performance networks are based on Ethernet technology but also use IP to allow them to scale to sizes beyond what is possible with Ethernet alone.

In addition to IP, many audio protocols borrow other components from the IT world including a basic means to transport audio data, network quality of service (QoS), synchronization, stream description, and connection management. The principal differences between the different current-generation IP audio protocols are in exactly how these components are used.

### AES67

This AES standard was developed with the primary goal of enabling interoperability between numerous audio network protocols that, despite all being IP based, were unable to communicate with each other. The standard specifies PTPv2 for synchronization and defines a "media profile" for PTP that is compatible with the PTP synchronization system used by SMPTE ST 2110. AES67 uses RTP/IP for transport and defines some common signaling and stream description mechanisms to encourage interoperation. AES67 was published in 2013.

### Ravenna

Ravenna was developed by Lawo division ALC NetworX in 2012 for use in Lawo products and for low-cost licensing to other, mainly European, audio manufacturers. Ravenna offers extensive configuration options with respect to QoS and PTP configuration. SDP, RSTP, and other standard protocols used for connection management functions. PTPv2 is used for synchronization with an emphasis on precision using PTP support in network equipment. A profile concept makes Ravenna a very flexible technology. One of the Ravenna profiles is AES67.

## Q-LAN

Q-LAN is the native IP protocol used by QSC to create networked audio connections between the various components of their Q-Sys product line. Q-LAN was included in the original Q-Sys launch in 2009. Because all system software in a Q-Sys system is updated together, QSC has been able to update Q-LAN over time. Q-LAN now uses the same PTPv2 synchronization and RTP/IP audio transport as AES67. Several configuration options are available with respect to QoS. Connection management is achieved using proprietary QSC protocols in concert with the Q-Sys Designer user interface and system configuration software. In addition to its native Q-LAN protocol, Q-Sys also features support for Dante, AES67, and several flavors of VoIP.

## Dante

Introduced in 2006 by Audinate, an Australian research firm, Dante forgoes RTP and uses simple UDP/IP messaging to transport audio data. IP's DiffServ system with fixed DSCP values is used for QoS. Synchronization is achieved using the original PTPv1. Connection management and stream descriptions are all handled by proprietary protocols aided by Apple's Bonjour service discovery system.

With recent firmware, Dante systems also support basic AES67. Because AES67 uses PTPv2 for synchronization, some care must be taken because both PTPv1 and PTPv2 must be running on the same network and synchronized with each other.

The Dante ecosystem is robust in terms of the number of different products that support the protocol, the number of audio engineers trained in how to install and use it, and the software, supplied by Audinate, used to configure and monitor a Dante system.

## VoIP and RTP

Voice over IP (VoIP) might be considered the original IP audio protocol. Most of our telephone calls today use these protocols at least at some point in their path. The core VoIP standard, RFC 3550, defining the Real-time Transport Protocol (RTP), was published in January 1996. Many of our professional IP audio protocols use RTP making them, essentially, very fancy VoIP systems. VoIP applications have relaxed synchronization requirements compared to professional audio applications and so usually do not use a separate precision synchronization component and instead rely on data arrival times to maintain synchronization.

## Sidebar: Synchronization Protocols

### Precision Time Protocol

Precision Time Protocol (PTP) was originally developed for scientific and industrial applications. The protocol delivers highly accurate time which can, in turn, be used to synchronize multiple devices on a local area network. PTP works reasonably well on a standard Ethernet network, delivering time with error in the range of microseconds. By using PTP-aware network equipment, accuracy is improved to less than 1 nanosecond.

Time is distributed from a designated grandmaster device on the network to any number of slaves. Time information is delivered in an IPv4 multicast *Sync* message sent at regular intervals from the grandmaster. To assure high accuracy, the time required to deliver these messages to slaves individually is measured and compensated for. This measurement is done with separate periodic *Delay\_request/response* message exchange with the grandmaster initiated individually by the slaves.

The grandmaster on a network is selected by a feature of PTP known as the best master clock algorithm (BMCA). Using an *Announce* message, each PTP device broadcasts information about its clocking abilities. The device with the best clock is elected grandmaster. In the event of a grandmaster failure, a new election is invoked, and the winner assumes the grandmaster role replacing the failed unit.

#### IEEE 1588-2002 (PTPv1)

The original PTP standard was released in 2002. PTPv1 includes the concept of a *boundary clock*. Network equipment containing a boundary clock is said to be *PTP aware*. The boundary clock synchronizes to a directly connected grandmaster or another boundary clock and then serves as a master for any directly connected slaves. The use of boundary clocks limits the need to pass PTP messages through the network where they may be delayed in network equipment that is not aware that PTP messages are time critical.

Although PTPv1 is still widely supported, most applications have now transitioned to PTPv2. An exception is the native Dante audio protocol.

#### IEEE 1588-2008 (PTPv2)

The updated PTP version released in 2008 is architecturally similar to the 2002 version but is functionally incompatible. It is possible to run PTPv1 and PTPv2 on the same network though they will act as separate clocks. Any PTP-aware equipment is typically configured for either PTPv1 or PTPv2 but cannot do both and, depending on how the implementation is designed, may or may not prevent the other protocol variant from working.

# Dropbox Integrated AV Production Environment

PTPv2 contains several new options including a transparent clock mode that improves the time required for slaves to initially synchronize. A native Ethernet protocol variant is included and has been adopted for use on AVB networks. The concept of a *profile* is introduced which allows the protocol to be tuned for specific applications. Profiles have subsequently been published for telecommunications, electrical power distribution, financial markets, and media applications.

## IEEE 1588-2019 (PTPv2.1)

A recent update is backward compatible with PTPv2 and introduces several new features that improve accuracy and robustness. These improvements have not yet been put to use in media systems. Devices supporting this new standard should be compatible with the PTPv2 devices, systems, and profiles currently deployed.

## SMPTE ST 2059

The SMPTE ST 2059 standard comes in two parts, 2059-1 and 2059-2. 2059-2 describes the use of PTPv2 for synchronization in media applications. For this purpose, it defines a PTP profile specifying operating and timing parameters. 2059-1 describes how to derive legacy synchronization signals such as black burst and tri-level sync from PTP time.

## AES67 media profile

Like SMPTE ST 2059, AES67 defines a PTP profile for media applications. The two profiles are largely compatible, and the AES has published recommendations in AES-R16 for configuring PTP in a way that is interoperable with both standards.

## Network Time Protocol

Network Time Protocol (NTP) has been in operation since the 1980s and is the protocol used by most computers to automatically set their clocks. Clock accuracy using NTP on small networks is on the order of 1 millisecond. Although NTP is less accurate than PTP, it is significantly more scalable providing synchronization services for large enterprises and even the Internet.

Where PTP uses a master-slave architecture with all slaves synchronizing to a single authoritative grandmaster, NTP is able to combine readings from multiple clocks to produce a robust composite time for synchronization.